


HERETICAL RESILIENCE (TO REPAIR IS HUMAN)

Ryn Daniels - @rynychantress
QCon New York 2018



Bridget Kromhout ✓

@bridgetkromhout

Following



Shoutout by @jezhumble at #srecon to @rynchantress for when they were awarded Etsy's three-armed sweater by @allspaw - we should all learn from failure. 100



12:19 PM - 28 Mar 2018 from Agnew, Santa Clara




MY SIDE OF THE STORY

AKA: A DRAMATIC
RETELLING OF THE
TIME I NEARLY BROKE
ETSY DOT COM

















































@RYNCHANTRESS


QCON NYC 2018










+

+

=



pod







Ryn Daniels

@rynchantress



So pleased to have won this year's Three-Armed Sweater Award for that time I accidentally upgraded apache! Thanks @allspaw! <3 #opslife



THE POST-MORTEM

AKA: WHAT THE HECK ACTUALLY JUST HAPPENED?


THE POST-MORTEM


AKA: WHAT THE HECK ACTUALLY JUST HAPPENED?

AKA: WHAT DID WE LEARN?



HOW DID THE
SITE STAY UP?





LESSON 1

Always keep 7 servers out of config management, just in case.

LESSON 1


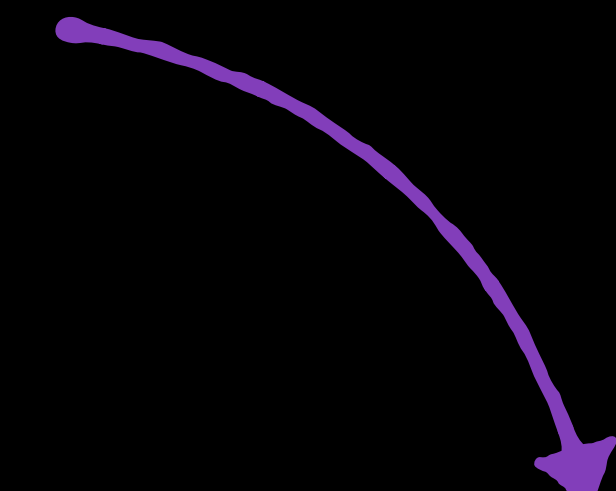

CONSIDER FALLBACKS FOR AUTOMATION

DISTRUSTING YOUR AUTOMATION

- How will you detect problems?
- How easily can you test your automation?
- Can you turn the automation off?
- Do you remember how to do the thing manually?



HOW DID WE
RESPOND SO
FAST?





LESSON 2

Create a Slack Team in charge of maintaining a proper amount of slack in case of incidents.

LESSON 2

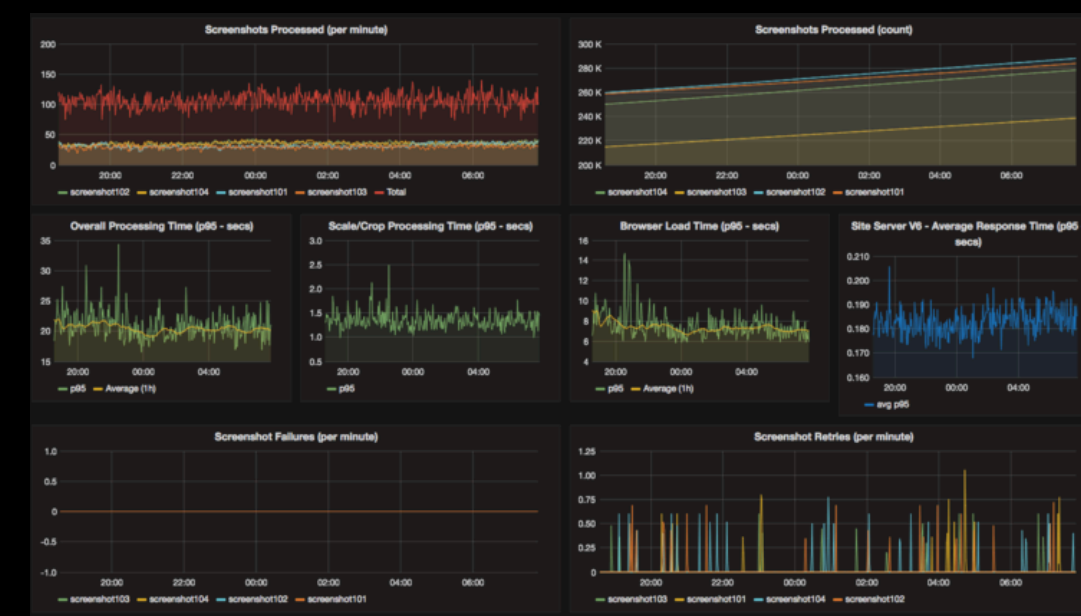
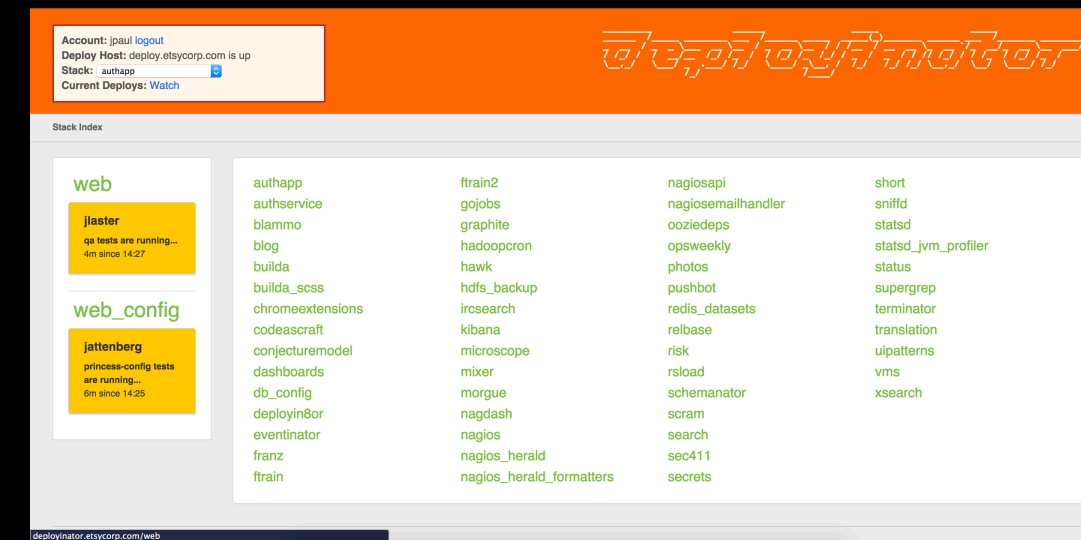
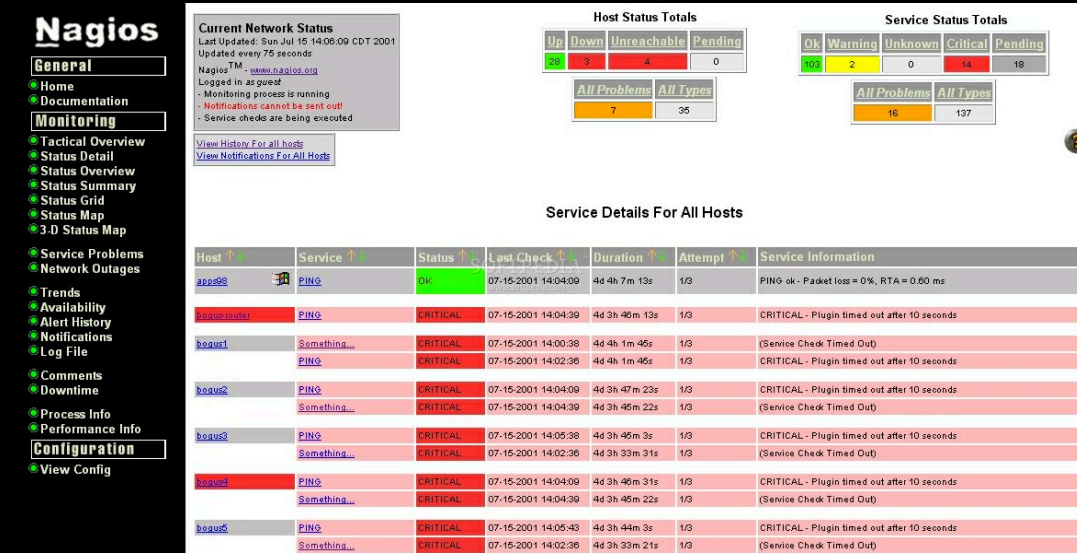
MAINTAIN ADAPTIVE CAPACITY

TWIDDLING YOUR THUMBS

- How do people ask each other for help?
- Which teams have more or less slack?
- What happens after work gets rearranged?



WHAT COULDN'T
WE SEE?





Nagios

Current Network Status
Last updated: Sun Jul 15 14:02:09 CDT 2011
Updated every 75 seconds
Nagios: ryan@lab101.us
Logged in as ryan
Monitoring process is running
No configuration errors to report
Service alerts are being avoided

Host Status Totals
Up: 2, Down: 0, Unreachable: 0, Pending: 0
All Problems: 0, All Types: 0

Service Status Totals
Up: 0, Warning: 0, Unknown: 0, Critical: 127
All Problems: 127, All Types: 127

Service Details For All Hosts

Host	Service	Status	Last Check	Duration	Attempts	Service Information
lab101	ping	UP	07-15-2011 14:04:09	44.3s 46m 13s	0/3	PING ok - Packet loss = 0%, RTT = 0.00 ms
lab101	ssh	CRITICAL	07-15-2011 14:04:09	44.4s 1m 46s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:02:36	44.4s 1m 46s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:04:09	44.3s 47m 23s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:04:09	44.3s 46m 23s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:02:36	44.3s 32m 31s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:04:09	44.3s 46m 31s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:04:09	44.3s 46m 22s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:06:43	44.3s 46m 3s	0/3	CRITICAL: Plugin timed out after 10 seconds
lab101	ssh	CRITICAL	07-15-2011 14:02:36	44.3s 32m 21s	0/3	CRITICAL: Plugin timed out after 10 seconds

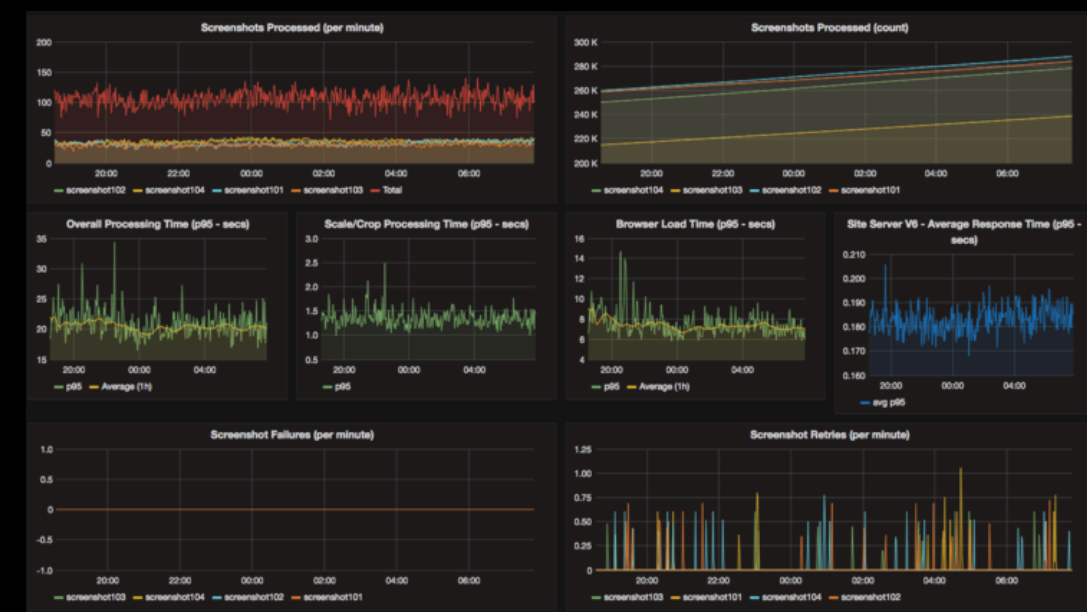



Deployinator

Account: jwald logged
Deploy Host: deploy.attyscorp.com is up
Stack: webapp
Current Deploys: Watch

Stack Index

web	authapp	ftrain2	nagiosapi	short
blaster	authservice	gojobs	nagiosemailhandler	sniffd
blammo	blammo	graphite	ooziedeps	statsd
blog	blog	hadoopcron	opieweeky	statsd_jvm_profiler
bulldo	bulldo	hawk	photos	status
bulldo_scs	bulldo_scs	hdfs_backup	puathbot	supergrsp
chromeextensions	chromeextensions	ircsearch	redis_datasets	terminator
codecraft	codecraft	kibana	rebase	translation
conjecturemodel	conjecturemodel	microscope	risk	upatterns
dashboards	dashboards	mixer	rsload	vms
db_config	db_config	morigue	schemanator	xsearch
deployinator	deployinator	nagiosdash	scram	
eventinator	eventinator	nagios	search	
franz	franz	nagios_herald	sec411	
ftrain	ftrain	nagios_herald_formatters	secrets	





Nagios

Current Network Status
Last Update: Sun Jul 15 14:06:09 CDT 2001
123 Active Alerts: 75 Alerts
Nagios™ v3.5.4-2001.02
Logged in as user@...
Monitoring process is running
Notifications cannot be sent out
Service checks are being executed

Host Status Totals			
Up	Down	Unreachable	Pending
10	0	0	0
All Problems: All Types: 0			

Service Status Totals			
OK	Warning	Unknown	Critical
16	2	0	18
All Problems: All Types: 18			

Service Details For All Hosts

Host	Service	Status	Last Check	Duration	Attempts	Service Information
server1	ping	OK	07-15-2001 14:04:09	44.46 7m 13s	1/3	FINO ok: Packet loss = 0%, RTT = 0.80 ms
server1	ssh	CRITICAL	07-15-2001 14:04:30	44.34 45m 15s	1/3	CRITICAL: Plugin timed out after 10 seconds
server1	snmp	CRITICAL	07-15-2001 14:02:36	44.46 1m 45s	1/3	CRITICAL: Plugin timed out after 10 seconds
server1	snmp	CRITICAL	07-15-2001 14:02:36	44.46 1m 45s	1/3	CRITICAL: Plugin timed out after 10 seconds
server2	ping	OK	07-15-2001 14:04:09	44.34 47m 22s	1/3	CRITICAL: Plugin timed out after 10 seconds
server2	ssh	CRITICAL	07-15-2001 14:04:30	44.34 45m 22s	1/3	CRITICAL: Plugin timed out after 10 seconds
server2	snmp	CRITICAL	07-15-2001 14:04:30	44.34 45m 22s	1/3	CRITICAL: Plugin timed out after 10 seconds
server2	snmp	CRITICAL	07-15-2001 14:04:30	44.34 45m 22s	1/3	CRITICAL: Plugin timed out after 10 seconds
server3	ping	OK	07-15-2001 14:04:09	44.34 46m 31s	1/3	CRITICAL: Plugin timed out after 10 seconds
server3	ssh	CRITICAL	07-15-2001 14:04:30	44.34 33m 31s	1/3	CRITICAL: Plugin timed out after 10 seconds
server3	snmp	CRITICAL	07-15-2001 14:04:30	44.34 33m 31s	1/3	CRITICAL: Plugin timed out after 10 seconds
server3	snmp	CRITICAL	07-15-2001 14:04:30	44.34 33m 31s	1/3	CRITICAL: Plugin timed out after 10 seconds
server4	ping	OK	07-15-2001 14:04:42	44.34 46m 31s	1/3	CRITICAL: Plugin timed out after 10 seconds
server4	ssh	CRITICAL	07-15-2001 14:04:30	44.34 33m 21s	1/3	CRITICAL: Plugin timed out after 10 seconds
server4	snmp	CRITICAL	07-15-2001 14:02:36	44.34 33m 21s	1/3	CRITICAL: Plugin timed out after 10 seconds

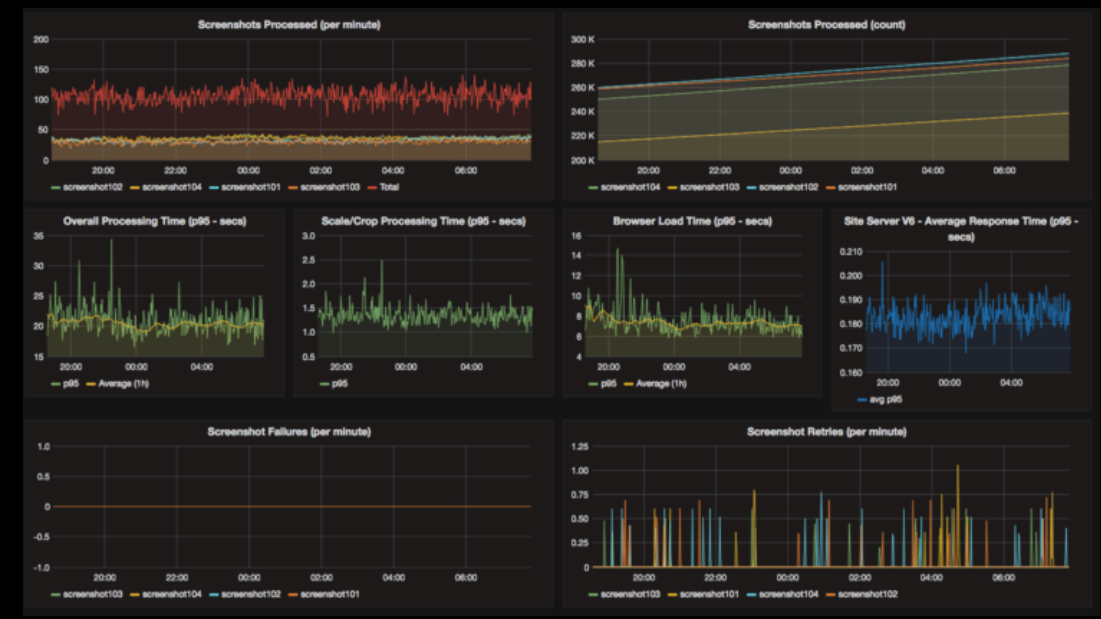


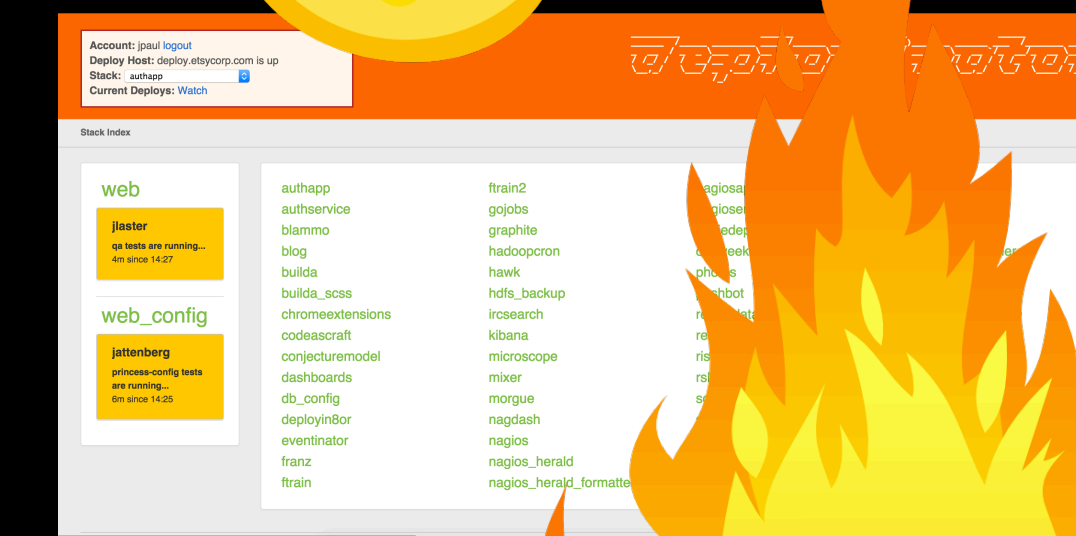
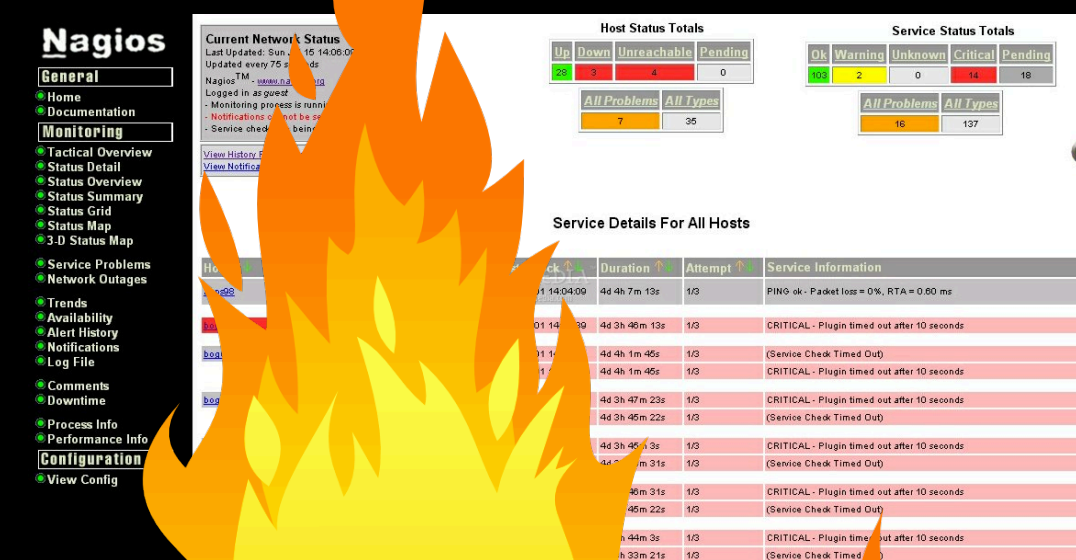
Deployinator

Account: jared login
Deploy Host: deploy.deployinator.com is up
Back: admin
Current Deploy: Watch

Stack Index

Category	Item	Item	Item	Item
web	authapp	train2	nagiosapi	short
	authservice	gojobs	nagiosmailhandler	sniff
	blammo	graphite	nagiosdeps	statusd
	blog	hadoopcron	opsweekly	statusd_ym_profilor
	builda	hawk	photos	status
	builda_scss	hdfs_backup	pushbot	superprep
	chromextensions	lrsearch	redis_datasets	terminator
	codecraft	klbina	rebase	translation
	conjecturemodal	microscope	risk	uploadstama
	dashboards	mixer	rload	vms
db_config	morgue	schemanator	xsearch	
deployin8or	nagdash	scram		
eventinator	nagios	search		
franz	nagios_herakd	sec411		
frain	nagios_herakd_formatters	secrets		





LESSON 3

Buy a couple botnets to DDoS your monitoring tools every now and then.

LESSON 3

UNDERSTAND THE DEPENDENCIES
IN YOUR TOOLING

WATCHING THE WORLD BURN

- What do your monitoring/automation/orchestration tools depend on?
- Who watches the watchers?
- How do you communicate internally and externally?
- Do you have backup tools?



WHAT
ACTUALLY WENT
WRONG WITH
CHEF?



LESSON 4

Always label your dragons.

LESSON 4

MAKE INFORMED DECISIONS ABOUT
WHICH YAKS TO SHAVE.

CHOOSING YOUR YAKS WISELY

- Which teams have sufficient slack?
- Can a problem be avoided if not solved?
- What are the tradeoffs and opportunity costs?
- Who has the precision yak razors?



WHO DIGS INTO
THE WEIRD
THINGS?

LESSON 4.5

Hire the person who created the primary language your site is written in.
(This always scales.)

LESSON 4.5

DEVELOP DEPTH OF INTER-TEAM RELATIONSHIPS

FINDING YOUR OWN RASMUS

- Which areas only have one (or two) people who understand them?
- How is information shared within your organization?
- What behaviors are rewarded?



STELLA

Report from the SNAFUcatchers Workshop on Coping With Complexity

Brooklyn NY, March 14-16, 2017

- 1. tl;dr and Executive Summary
 - 1.1 tl;dr
 - 1.2 Executive Summary
- 2. Introduction
 - 2.1 About the SNAFUcatchers consortium and the STELLA meeting
 - 2.2 The focus on handling anomalies
 - 2.3 The above-the-line/below-the-line framework
- 3. Cases
 - 3.1 Catching the Apache SNAFU
 - 3.2 Catching the Travis CI SNAFU
 - 3.3 Catching the Logstash SNAFU
 - 3.4 Observations on the cases
 - 3.4.1 Features of the anomalies
 - 3.4.2 Features of the anomaly responses
- 3.5 Observations on the postmortem process
- 4. Themes
 - 4.1 Capturing the value of anomalies through postmortems
 - 4.1.1. Technical issues in postmortems:
 - 4.1.2. Social issues in postmortems
 - 4.2 Blame versus sanction in the aftermath of anomalies
 - 4.3 Controlling the costs of coordination during anomaly response
 - 4.3.1 Offloading work to low-tempo periods
 - 4.3.2 Providing expertise on demand
 - 4.3.3 Supporting communication and coordination with tools
 - 4.4 Supporting anomaly response through improved visualizations



Winter storm STELLA

Woods' Theorem: *As the complexity of a system increases, the accuracy of any single agent's own model of that system decreases rapidly.*

1. tl;dr and Executive Summary

WHAT HAPPENED AFTERWARDS?



laurensperber
Etsy HQ



LESSON 5

Give people ill-fitting clothing
when they mess up.

LESSON 5

ENCOURAGE ORGANIZATIONAL LEARNING

A WARNING TO OTHERS

- How do people respond to incidents?
- What happens after an incident?
- How are remediation items prioritized?
- What happen to the bandaid solutions?



TECHNOLOGY CAN BE ROBUST.*
ONLY HUMANS CAN BE RESILIENT.

*for some already-known, pre-defined subset of problems



John Allspaw

@allspaw

Following



If the story is only about software and/or hardware, it cannot be about "resilience" - it's likely about "robustness."

Resilience is about *people* and the explicit support for their ability to adapt to unforeseen circumstances.

7:23 PM - 26 Jun 2018

1. UNDERSTAND YOUR AUTOMATION
2. MAINTAIN ADAPTIVE CAPACITY
3. KNOW YOUR DEPENDENCIES
4. BUILD CROSS-TEAM RELATIONSHIPS
5. ALWAYS BE LEARNING

1. UNDERSTAND YOUR AUTOMATION
2. MAINTAIN ADAPTIVE CAPACITY
3. KNOW YOUR DEPENDENCIES
4. BUILD CROSS-TEAM RELATIONSHIPS
5. ALWAYS BE LEARNING





THANK YOU!